III.B.1. The ONCOCIN Project

The central goal of the ONCOCIN project, under Professor Ted Shortliffe at Stanford Medical School, is to develop programs that can provide advice in oncology protocol management, similar in quality to that given by human experts. The design goals are also to ensure that the system is easy to use and acceptable to physicians. The underlying artificial intelligence research of this project seeks to improve the interactive process, both for the developer of a knowledge-based system, and for the intended end user. In addition, we have emphasized clinical implementation of the developing tool so that we can ascertain the effectiveness of the program's interactive capabilities when it is used by physicians who are caring for patients and are not involved in the computer-based research activity.

ONCOCIN is a large interdisciplinary effort that has involved over 35 individuals since the project's inception in July 1979. The work is currently in its seventh year and significant progress has been made this past year. The OPAL knowledge acquisition system became operational and several new oncology protocols have been entered using this system. As anticipated, we have increased the speed and ease with which protocols can be added to the ONCOCIN knowledge base. Using the protocols entered through OPAL, we have begun experimental testing of the workstation version of ONCOCIN in the Stanford oncology clinic.

We have connected the various parts of the system, and have demonstrated that we have the capability to run ONCOCIN with the reasoning program and interface program on different machines in a communications network. The current version of the program is running on a single workstation, but future versions may take advantage of the multiple machine option. To increase the speed at which we are able to test protocols entered into ONCOCIN, we have developed additional programs to test real and synthetic cases without user interaction; these are then reviewed by our collaborating clinicians.

We have also developed a workstation-based program, OPUS, to help clinicians determine which protocols are appropriate for specific patients. We have been using it in the clinic setting since the end of 1985. Thus, in addition to providing an information resource about protocols, the use of a graphically-oriented program provided a way to learn about the software style and hardware used in the workstation version of ONCOCIN.

The new workstation version of ONCOCIN is still under evaluation. The very promising results of evaluating the older mainframe version was documented in two evaluation papers that appeared in clinical journals[1]. As a further step in our dissemination plans, we are planning experimental installation of ONCOCIN workstations in private oncology offices in San Jose and San Francisco. An application proposing this project is currently under review.

We have also continued our basic research on the design of advanced therapy-planning programs, the ONYX project. We have developed a model for planning which includes techniques from the fields of artificial intelligence, simulation, and decision analysis. Our goal is to allow ONCOCIN to give advice in a wider range of situations; in particular, the system should be able to recommend plans for patients who have an unusual response to chemotherapy.

[1]Hickam, D.H., Shortliffe, E.H., Bischoff, M.B., Scott, A.C., Jacobs, C.D. A study of the treatment advice of a computer-based cancer chemotherapy protocol advisor (Memo KSL-85-21). Annals of Internal Medicine 103(6 pt 1):928-936 (1985).

A great deal of interest in ONCOCIN has been shown by the medical, computer science, and lay communities. We are frequently asked to demonstrate the program to Stanford visitors (both the prototype system running in the clinic and the newer work transferring the system to professional workstations). We also demonstrated our workstation system in the Xerox exhibit in the trade show associated with AAAI-84 in Austin, Texas and IJCAI-85 in Los Angeles. Physicians have generally been enthusiastic about ONCOCIN's potential. The interest of the lay community is reflected in the frequent requests for magazine interviews and television coverage of the work. Articles about MYCIN and ONCOCIN have appeared in such diverse publications as *Time* and *Fortune*, and ONCOCIN has been featured on the "NBC Nightly News," the PBS "Health Notes" series, and "The MacNeil-Lehrer Report." Due to the frequent requests for ONCOCIN demonstrations, we have produced a videotape about the ONCOCIN research which includes demonstrations of our professional workstation research projects and the 2020-based clinic system. The tape has been shown at several national meetings, including the 1984 Workshop on Artificial Intelligence in Medicine, the 1984 meeting of the Society for Medical Decision Making, and the 1985 meeting of the Society for Research and Education in Primary Care Internal Medicine. The tape has also been shown to both national and international researchers in biomedical computing. We have also completed an updated tape of our activities for demonstration purposes.

III.B.2. The Internist-I Project

The principal objective of Internist-I project, under Professor Jack Myers and Dr. Randolph Miller at the University of Pittsburgh, is to develop a high-level computer diagnostic program in the broad field of internal medicine as an aid in the solution of complex diagnostic problems. To be effective, the program must be capable of multiple diagnoses (related or independent) in a given patient.

A major achievement of this research undertaking has been the design of a program called INTERNIST-I, along with an extensive medical knowledge base. This program has been used over the past decade to analyze many hundreds of difficult diagnostic problems in the field of internal medicine. These problem cases have included cases published in medical journals (particularly Case Records of the Massachusetts General Hospital, in the New England Journal of Medicine), CPCs, and unusual problems of patients in our Medical Center. In most instances, but by no means all, INTERNIST-I has performed at the level of the skilled internist, but this experience has also highlighted several areas of the program needing further improvement.

The development of the QUICK MEDICAL REFERENCE (QMR) system, under the leadership of Dr. Miller, has allowed us to distribute the INTERNIST-I knowledge base in a modified format to over 20 other academic medical institutions. The entire QMR program runs on individual IBM PC-AT workstations. The knowledge base can thereby be used as an "electronic textbook" in medical education at all levels -- by medical students, residents and fellows, and faculty and staff physicians. This distribution is continuing to expand.

The INTERNIST-I program has been used in recent years to develop patient management problems for the American College of Physician's Medical Knowledge Self-assessment Program, and to develop patient management problems and test cases for the Part III Examination and the developing computerized testing program of the National Board of Medical Examiners.

The project staff are continuing to expand the knowledge base of INTERNIST-I and to incorporate the diagnostic consultative program into QMR. The medical knowledge base has continued to grow both in the incorporation of new diseases and the modification of diseases already profiled so as to include recent advances in medical knowledge. Several dozen new diseases have been profiled during the past year. The computer program for the interrogative part of the diagnostic program is the main remaining task. An editor for the QMR knowledge base, as modified from the INTERNIST-I knowledge base, is nearing completion.

In the near future the project will be ready to incorporate into the QMR diagnostic consultant program the enhancements of the INTERNIST-I knowledge base, e.g. the use of "facets" of diseases or syndromes. This addition is expected to improve the performance of the diagnostic consultant program.

III.B.3. The PROTEAN Project

The goals of this project, under Professors Oleg Jardetzky and Bruce Buchanan at Stanford University, are related both to biochemistry and artificial intelligence: (a) use existing AI methods to aid in the determination of the 3-dimensional structure of proteins in solution (as opposed to crystallized proteins), and (b) use protein structure determination as a test problem for experiments with the AI problem solving structure known as the Blackboard Model. Empirical data from nuclear magnetic resonance (NMR) and other sources may provide enough constraints on structural descriptions to allow protein chemists to bypass the laborious methods of crystallizing a protein and using X-ray crystallography to determine its structure. This problem exhibits considerable complexity. Yet there is reason to believe that AI programs can be written that reason much as experts do to resolve these difficulties.

The molecular structure of proteins is essential for understanding many problems of medicine at the molecular level, such as the mechanisms of drug action. Using NMR data from proteins in solution will allow the study of proteins whose structure cannot be determined with other techniques, and will decrease the time needed for the determination.

Over the past year, we have constructed a prototype of such a program, called PROTEAN, designed on the blackboard model. It is implemented in BB1, a framework system for building blackboard systems that control their own problem-solving behavior. The reasoning component of PROTEAN directs the actions of the Geometry System (GS), a set of programs that performs the computationally intensive task of positioning portions of a molecule with respect to each other in three dimensions. Currently we have implemented two versions of the GS: an InterLISP version used for quickly testing ideas and developing prototypes of geometric routines on a LISP workstation; and a high performance version written in C and running in the UNIX environment, providing efficient computations on a VAX 11/780. We have coupled the reasoning and geometry programs with an IRIS graphics terminal (shared with SUMEX) which displays the evolving protein structures at several levels of detail. The display in three dimensions is essential to understanding the behavior of the reasoning and geometry systems, and provides essential insights on the problem solving process.

The current version of PROTEAN has five domain knowledge sources and five control knowledge sources that demonstrate the reasoning techniques described above for the assembly of a protein. Each domain knowledge source directs a small portion of the construction of the molecule. These knowledge sources develop partial solutions that position alpha helices, beta strands, and random coils at the Solid level and refine the resulting state families using all available distance constraints. The control knowledge sources determine which of the possible assembly actions is the best to perform at each stage of the problem solving.

The PROTEAN system [2] has been used to construct a complete solution at the solid level of detail for the Lac-repressor headpiece, a protein with 51 amino acids consisting of 4 random coil sections and three alpha helices[1] . In this work, the constraints were determined experimentally from NMR studies.

To demonstrate that our method is correct, we have applied PROTEAN to sperm whale myoglobin, a molecule whose crystal structure is known. By using distances between atoms in the crystal, distance constraints were applied to the eight helices in myoglobin

---

[1]Buchanan, B.G., Altman, A., Brinkley, J., Cornelius, C., Duncan, B., Hayes-Roth, B., Hewett, M., Lichtarge, O., Jardetzky, O.: *The Heuristic Refinement Method for Deriving Solution Structures of Proteins*. Report KSL-85-41. October 1985. In *Proceedings of the National Academy of Sciences*..

to determine if PROTEAN would reproduce the crystal structure. The family of solutions obtained from PROTEAN includes the actual structure of the molecule. Work is proceeding to include the heme group of myoglobin as a component and use constraints to other portions of the molecule to further restrict the state families obtained.

III.B.4. AIM Community Software Support

The SUMEX-AIM resource staff have made major efforts to assist other members of the SUMEX-AIM community in integrating the technologies needed for biomedical AI research. This is often achieved through direct contact with staff members at these institutions, at meetings and workshops, or via electronic mailing lists. For example, the SUMEX-AIM gateway software is well known for its versatility and reliability, and has been distributed to both the University of Chicago and Rutgers University through contacts made this way. The gateway running at Rutgers is quite similar to our own and uses Ethernet as its medium of transport. The gateway at the University of Chicago is a principle node on a fibre-optic network.

Dr. Charles Hedrick, Director of the Rutgers Computer Science Department computer facility, acknowledged our contribution to the Rutgers networking efforts as follows:

> **SUMEX Support of Networking at Rutgers**
>
> During the last two years, computing at Rutgers has become increasingly network-oriented. This change has included the Rutgers AIM project, as well as most other compute-intensive research areas within the University. Our primary computing resource now consists of workstations such as Suns or MicroVAXes for individual researchers, and super-minis at the departmental or project level. All of these systems are connected to each other via Ethernet. The workstations depend upon the super-minis for a number of services, including file storage, backup, and printing. In addition, there are more specialized services available via the Ethernet, including larger mainframes, a high-speed Xerox laser printer, and (shortly) the facilities of the John von Neumann Supercomputer Center. The Ethernet-based network has become critical to the operation of most of our computing facilities.
>
> To a large extent the network technology now used at Rutgers is based on work done at Stanford, and particularly at SUMEX. The staff at SUMEX have given us significant assistance in importing this technology to Rutgers. Much of the AI research at Rutgers is done using Xerox Interlisp-D workstations. These workstations require network file service via either the PUP or XNS protocols. Our primary file server is a Pyramid Unix system that uses the Unix PUP implementation developed at SUMEX.
>
> During the last year, the amount of Ethernet work at Rutgers became large enough that a single Ethernet was no longer sufficient. At that point we created 4 separate Ethernets. They are connected with an Ethernet gateway developed at SUMEX. SUMEX personnel supplied us with a description of the hardware necessary. It could be assembled from commercially-available pieces. They also supplied us with the software to be run in the gateway, and the development tools necessary to compile and down-load it. Although our setup is very similar to the one at Stanford, a certain amount of customization was necessary for the environment here. Most of this work was done by the SUMEX staff. They also provided assistance for the rest. We have the highest praise for the quality of this software, and for the support that we have gotten from the SUMEX staff, particularly Bill Yeager. This is the only piece of software we have that has never crashed or otherwise misbehaved. Its overall quality, including such items as administrative tools, is better than most commercially-supported products we have seen.
>
> We consider this sort of research and technology transfer to be an important benefit of SUMEX. The tendency now is for each research group to have its own equipment and its own support staff. This is both inevitable and useful. However it is important for the AIM community to have an organization such as

SUMEX that can investigate new computing technology and make it available to the rest of the community.

III.B.5. Remote Virtual Graphics

Lisp workstations of various types have proven extremely powerful, both as development environments for artificial intelligence research and as vehicles for disseminating AI systems into user communities. In addition to the compact, inexpensive computing resources workstations provide, high-quality graphics play a key role in their power. Such graphics systems have become indispensable for understanding the complex data structures involved in developing and debugging large AI systems and are important in facilitating user access to working programs (e.g., for ONCOCIN and PROTEAN). However, as we move towards a distributed workstation computing environment for AI research in the SUMEX-AIM community (and move away from the centralized, shared DEC 2060), a number of technical obstacles must be overcome. One of the most important is to eliminate the need for the workstation display to be situated close to the workstation computing engine.

In the past, members of the SUMEX-AIM community have often watched each others programs work by linking their CRT terminals to the text output of a running program on the SUMEX 2060. In the case of workstations, though, it is much more difficult to link across several networks to view the complex graphics output of a program. Even locally, it is important to make graphical interaction with workstations across campus or from home possible. One would like to be able to provide the same powerful graphical tools and programming environment that are available to a user sitting in front of the workstation to the remote user if that user has a low cost bit-mapped display and mouse. In order to accomplish this, it is necessary to capture and encode the many graphics operations involved so that they can be sent over a relative low-speed network connection with the same interactive facility as if one had the display connected through the dedicated high-speed (30 Mhz) native vendor display/workstation connection.

In order to study the feasibility of remote access to workstations to satisfy these requirements, we have been experimenting with virtual graphics protocols (VGP) to capture graphics constructs and interactions at a high level. Such a protocol, the Virtual Graphics Terminal Service (VGTS), was developed at Stanford in the Computer Science distributed systems group[1]. The VGTS provides tools to define objects like windows, lines, rectangles, circles, bitmaps, ellipses, splines, and graphics events like mouse clicks independently of the graphics hardware and operating systems. This encoding minimizes the communication bandwidth required between cooperating hosts, for example, to remotely draw a line.

Over the past year, an implementation of this protocol was developed and installed in the operating system of a Xerox 1186 Lisp workstation so that its presence would be transparent to the programmer. This means that if one connects to such a LISP workstation from a SUN workstation (running suitable VGTS software), the Lisp machine graphics will be sent over the net and reconstructed on the SUN workstation without changes to the application program running. This implementation has worked very well in early experiments so that over an Ethernet, the remote response time is quite close to the response time on the Lisp machine itself.

As a consequence of this work, we have demonstrated the feasibility of remotely using LISP workstations over an Ethernet to take advantage of their graphics programming environment. Work in the virtual graphics group at Stanford also suggests that the VGTS can be used effectively over the ARPANET (56 KBits/sec), at least for some kinds of applications, with excellent response time. Much more work needs to be done

---

[1]see Keith A. Lantz, David R. Cheriton, and William I. Nowicki, *Third Generation Graphics for Distributed Systems*, Computer Systems Laboratory Report, Stanford University, December 21, 1982.

in this area to fully develop and integrate these capabilities into Lisp machine systems and to insure that cross-country connections will indeed give usable response time. Success of this work will mean that one can use LISP machine systems from TYMNET, UNINET, ARPANET, or an Ether TIP connection throughout the SUMEX-AIM community.

III.C. Administrative Changes

There have been few administrative changes within the project this past year other than some turn-over in staff. In the coming year, Professor Shortliffe will be on sabbatical at the University of Pennsylvania. During this time, Professor Feigenbaum will reassume the role of principal investigator.

More substantial developments are expected in the coming year as well. First we will be moving the Medical Computer Science and SUMEX offices into newly constructed space within a Stanford Medical School Office Building. We will occupy approximately 6500 square feet which almost doubles the space available to us. This space is designed to improve the interactions within our groups.

The other development will be the design of a cost recovery system as we phase out BRTP subsidy of the DEC 2060 facility. We plan to use the continuing component of NIH support to retain our no fee for service policy for non-Stanford projects for as long as possible. We will establish a cost center to allocate the uncovered expenses to Stanford user projects.

III.D. Resource Management and Allocation

Early in the design of the SUMEX-AIM resource, an effective management plan was worked out with the Biotechnology Resources Program (now Biomedical Research Technology Program) at NIH to assure fair administration of the resource for both Stanford and national users and to provide a framework for recruitment and development of a scientifically meritorious community of application projects. This structure has been described in some detail in earlier reports and is documented in our recent renewal application. It has continued to function effectively as summarized below.

- The AIM Executive Committee meets regularly by teleconference to advise on new project applications, discuss resource management policies, plan workshop activities, and conduct other community business. The Advisory Group meets together at the annual AIM workshop to discuss general resource business and individual members are contacted much more frequently to review project applications. (See Appendix A on page 187 for a current listing of AIM committee membership).

- We have actively recruited new application projects and disseminated information about the resource. The number of formal projects in the SUMEX-AIM community still runs at the capacity of our computing resources. With the development of more decentralized computing resources within the AIM community outside of Stanford (see below), the center of mass of our community has naturally shifted toward the growing number of Stanford applications and core research projects. We still, however, actively support new applications in the national community where these are not able to gain access to suitable computing resources on their own.

- With the advice of the Executive Committee, we have awarded pilot project status to promising new application projects and investigators and where appropriate, offered guidance for the more effective formulation of research plans and for the establishment of research collaborations between biomedical and computer science investigators.

- We have carefully reviewed on-going projects with our management committees to maintain a high scientific quality and relevance to our biomedical AI goals and to maximize the resources available for newly developing applications projects. Several fully authorized and pilot projects have been encouraged to develop their own computing resources separate from SUMEX or have been phased off of SUMEX as a result and more productive collaborative ties established for others.

- We have continued to provide active support for the AIM workshops. The last one was held in Washington, DC, hosted by the National Library of Medicine under Drs. Lindberg and Kingsland.

- We will have to partly abandon our policy of no fee-for-service for projects using the SUMEX resource. This policy had effectively eliminated the serious administrative barriers that would have blocked our research goals of broader scientific collaborations and interchange on a national scale within the selected AIM community. We will attempt to minimize these barriers for national projects but will be obliged to recover costs not supported by BRTP as we phase out support for our mainframe computers (DEC 2060, 2020, and VAX 11/780).

- We have tailored resource policies to aid users whenever possible within our research mandate and available facilities. Our approach to system scheduling, overload control, file space management, etc. all attempt to give users the greatest latitude possible to pursue their research goals consistent with fairly meeting our responsibilities in administering SUMEX as a national resource.

As indicated above, we have sought to retain SUMEX resources for new projects, those exploring new areas in biomedical AI applications and those in such an early state of feasibility that they are unable to afford their own computing resources. This policy has worked effectively as seen from the following lists of terminated projects and projects now using their own computing resources at other sites:

Projects Moved All or In Part to Other Machines:

Stanford Projects:

- GENET [Brutlag, Kedes, Friedland - IntelliCorp]

National Projects:

- Acquisition of Cognitive Procedures (ACT) [Anderson - CMU]
- Chemical Synthesis [Wipke - UC Santa Cruz]
- Simulation of Cognitive Processes [Lesgold - Pittsburgh]
- PUFF [Osborne, Feigenbaum, Fagan - Pacific Medical Center]
- CADUCEUS/INTERNIST [Pople, Myers - Pittsburgh]
- Rutgers [Amarel, Kulikowski, Weiss - Rutgers]
- MDX [Chandrasekaran - Ohio State]
- SOLVER [P. Johnson - University of Minnesota]

Completed Projects Summary

Stanford Projects:

- DENDRAL [Lederberg, Djerassi, Buchanan, Feigenbaum]
- MYCIN [Shortliffe, Buchanan]
- EMYCIN [Shortliffe, Buchanan]
- CRYSALIS [Feigenbaum, Engelmore]
- MOLGEN I [Feigenbaum, Brutlag, Kedes, Friedland]
- AI Handbook [Feigenbaum, Barr, Cohen]
- AGE Development [Feigenbaum, Nii]

National Projects:

- Ventilator Management [Osborne, Feigenbaum, Fagan - Pacific Medical Center]

- Higher Mental Functions [Colby - USC]

III.E. Dissemination of Resource Information

Throughout the history of the SUMEX-AIM resource, we have made extensive efforts at disseminating the AI technology developed here. This has taken the form of many publications -- over 45 combined books and papers are published per year from the KSL; wide distribution of our software including systems software and AI application and tool software, both to other research laboratories and for commercial development; production of films and video tapes depicting aspects of our work; and significant project efforts at studying the dissemination of individual applications systems such as the GENET community (DNA sequence analysis software) and the ONCOCIN resource-related research project (see 109).

Software Distribution

We have widely distributed both our system software and our AI tool software. We have no accurate records of the extent of distribution of the system codes because their distribution is not centralized and controlled. The recent programs such as TOPS-20 monitor enhancements, the Ethernet gateway and TIP programs, the SEAGATE AppleBus to Ethernet gateway, the PUP Leaf server, the SUMACC development system for Macintosh workstations, and our Lisp workstation programs are well-distributed throughout the ARPANET community and beyond.

We have made major efforts to assist other members of the SUMEX-AIM community in integrating the technologies needed for biomedical AI research. This is often achieved through direct contact with staff members at these institutions at meetings and workshops or via electronic mailing lists. For example, the SUMEX-AIM gateway software is well known for its versatility and reliability, and has been distributed to both the University of Chicago and Rutgers University through contacts made this way. The gateway running at Rutgers is quite similar to our own and uses Ethernet as its medium of transport. The gateway at the University of Chicago is a principle node on a fibre-optic network. Dr. Charles Hedrick, Director of the Rutgers Computer Science Department facility, acknowledged our contribution to the Rutgers networking efforts as follows:

> **SUMEX Support of Networking at Rutgers**
>
> During the last two years, computing at Rutgers has become increasingly network-oriented. This change has included the Rutgers AIM project, as well as most other compute-intensive research areas within the University. Our primary computing resource now consists of workstations such as Suns or MicroVAXes for individual researchers, and super-minis at the departmental or project level. All of these systems are connected to each other via Ethernet. The workstations depend upon the super-minis for a number of services, including file storage, backup, and printing. In addition, there are more specialized services available via the Ethernet, including larger mainframes, a high-speed Xerox laser printer, and (shortly) the facilities of the John von Neumann Supercomputer Center. The Ethernet-based network has become critical to the operation of most of our computing facilities.
>
> To a large extent the network technology now used at Rutgers is based on work done at Stanford, and particularly at SUMEX. The staff at SUMEX have given us significant assistance in importing this technology to Rutgers. Much of the AI research at Rutgers is done using Xerox Interlisp-D workstations. These workstations require network file service via either the PUP or XNS protocols. Our primary file server is a Pyramid Unix system that uses the Unix PUP implementation developed at SUMEX.

During the last year, the amount of Ethernet work at Rutgers became large enough that a single Ethernet was no longer sufficient. At that point we created 4 separate Ethernets. They are connected with an Ethernet gateway developed at SUMEX. SUMEX personnel supplied us with a description of the hardware necessary. It could be assembled from commercially-available pieces. They also supplied us with the software to be run in the gateway, and the development tools necessary to compile and down-load it. Although our setup is very similar to the one at Stanford, a certain amount of customization was necessary for the environment here. Most of this work was done by the SUMEX staff. They also provided assistance for the rest. We have the highest praise for the quality of this software, and for the support that we have gotten from the SUMEX staff, particularly Bill Yeager. This is the only piece of software we have that has never crashed or otherwise misbehaved. Its overall quality, including such items as administrative tools, is better than most commercially-supported products we have seen.

We consider this sort of research and technology transfer to be an important benefit of SUMEX. The tendency now is for each research group to have its own equipment and its own support staff. This is both inevitable and useful. However it is important for the AIM community to have an organization such as SUMEX that can investigate new computing technology and make it available to the rest of the community.

We do have reasonably accurate records of the distribution of our AI tool software because the recipient community is more directly coupled to us and the distribution is centralized. Records indicate that over the past three years there have been about 1000 inquiries that have resulted in the distribution of written material about our software systems. It is likely that there have been a similar number of unrecorded or informal interactions on the part of the staff.

In the past year, we have distributed 73 copies of various software packages, including 18 EMYCIN systems, 41 MRS systems, 7 AGE systems, and 7 BB1 systems. In addition, a substantial amount of time was spent organizing the 106 files of the three volumes of the A/I Handbook so that it could be provided to two collaborators. And finally, we are organizing the DENDRAL software to make it available for distribution again. The CONGEN portion is essentially complete and being tested by an outside user and one of the original developers. Some progress has been made with the organizing of the GENOA portion. Several interested groups are awaiting its release.

Video Tapes and Films

The KSL and the ONCOCIN project have prepared several video tapes that provide an overview of the research and research methodologies underlying our work and that demonstrate the capabilities of particular systems. These tapes are available through our groups, the Fleischmann Learning Center at the Stanford Medical Center, and the Stanford Computer Forum and copies have been mailed to program offices of our various funding sponsors. The three tapes include:

- *Knowledge Engineering in the Heuristic Programming Project* -- This 20-minute film/tape illustrates key ideas in knowledge-based system design and implementation, using examples from ONCOCIN, PROTEAN, and knowledge-based VLSI design systems. It describes the research environment of the KSL and lays out the methodologies of our work and the long term research goals that guide it.

- *ONCOCIN Overview* -- This is a 30-minute tape providing an overview of the ONCOCIN project. It gives an historical context for the work, discusses the clinical problem and the setting in which the prototype system is being used, and outlines the plans for transferring the system to run on single-user workstations. Brief illustrations of the graphics capabilities of ONCOCIN on a Lisp workstation are also provided.

- *ONCOCIN Demonstration* -- This 1-hour tape provides detailed examples of the key components of the ONCOCIN system. It begins with a demonstration of the prototype system's performance on a time-shared mainframe computer and then shows each of the elements involved in transferring the system to Lisp workstations.

III.F. Suggestions and Comments

Resource Organization

We continue to believe that the Biomedical Research Technology Program is one of the most effective vehicles for developing and disseminating technological tools for biomedical research. The goals and methods of the program are well-designed to encourage building of the necessary multi-disciplinary groups and merging of the appropriate technological and medical disciplines.

Electronic Communications

SUMEX-AIM has pioneered in developing more effective methods for facilitating scientific communication. Whereas face-to-face contacts continue to play a key role, in the longer-term we feel that computer-based communications will become increasingly important to the NIH and the distributed resources of the biomedical community. We would like to see the BRTP take a more active role in promoting these tools within the NIH and its grantee community.

IV. Description of Scientific Subprojects

The following subsections report on the AIM community of projects and "pilot" efforts including local and national users of the SUMEX-AIM facility at Stanford. However, those projects admitted to the National AIM community which use the Rutgers-AIM resource as their home base are not explicitly reported here.

In addition to these detailed progress reports, abstracts for each project and its individual users are submitted on a separate Scientific Subproject Form. However, we have included here briefer summary abstracts of the fully-authorized projects in Appendix B on page 191.

The collaborative project reports and comments are the result of a solicitation for contributions sent to each of the project Principal Investigators requesting the following information:

I. SUMMARY OF RESEARCH PROGRAM
   A. Project rationale
   B. Medical relevance and collaboration
   C. Highlights of research progress
      --Accomplishments this past year
      --Research in progress
   D. List of relevant publications
   E. Funding support

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE
   A. Medical collaborations and program dissemination via SUMEX
   B. Sharing and interactions with other SUMEX-AIM projects
      (via computing facilities, workshops, personal contacts, etc.)
   C. Critique of resource management
      (community facilitation, computer services, communications services, capacity, etc.)

III. RESEARCH PLANS
   A. Project goals and plans
      --Near-term
      --Long-range
   B. Justification and requirements for continued SUMEX use
   C. Needs and plans for other computing resources beyond SUMEX-AIM
   D. Recommendations for future community and resource development

We believe that the reports of the individual projects speak for themselves as rationales for participation. In any case, the reports are recorded as submitted and are the responsibility of the indicated project leaders. The only exceptions are the respective lists of relevant publications which have been uniformly formatted for parallel reporting on the Scientific Subproject Form.

IV.A. Stanford Projects

The following group of projects is formally approved for access to the Stanford aliquot of the SUMEX-AIM resource. Their access is based on review by the Stanford Advisory Group and approval by Professor Feigenbaum as Principal Investigator.

In addition to the progress reports presented here, abstracts for each project and its individual users are submitted on a separate Scientific Subproject Form.

IV.A.1. GUIDON/NEOMYCIN Project

GUIDON/NEOMYCIN Project

William J. Clancey, Ph.D.
Department Computer Science
Stanford University

Bruce G. Buchanan, Ph.D.
Computer Science Department
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The GUIDON/NEOMYCIN Project is a research program devoted to the development of a knowledge-based tutoring system for application to medicine. This work derived from our first system, the MYCIN program. That research led to three sub-projects (EMYCIN, GUIDON, and ONCOCIN) described in previous annual reports. EMYCIN has been completed and its resources reallocated to other projects. GUIDON and ONCOCIN have become projects in their own right.

The key issue for the GUIDON/NEOMYCIN project is to develop a program that can provide advice similar in quality to that given by human experts, modeling how they structure their knowledge as well as their problem-solving procedures. The consultation program using this knowledge is called NEOMYCIN. NEOMYCIN's knowledge base, designed for use in a teaching application, is the subject material used by a family of instructional programs referred to collectively as GUIDON2. The problem-solving procedures are developed by running test cases through NEOMYCIN and comparing them to expert behavior. Also, we use NEOMYCIN as a test bed for the explanation capabilities incorporated in our instructional programs.

The purpose of the current contracts is to construct a knowledge-based tutoring system that teaches diagnostic strategies explicitly. By strategy, we mean plans for establishing a set of possible diagnoses, focusing on and confirming individual diagnoses, gathering data, and processing new data. The tutorial program has capabilities to recognize these plans, as well as to articulate strategies in explanations about how to do diagnosis. The strategies represented in the program, modeling techniques, and explanation techniques are wholly separate from the knowledge base, so that they can be used with many medical (and non-medical) domains. That is, the target program will be able to be tested with other knowledge bases, using system-building tools that we provide.

B. Medical Relevance and Collaboration

There is a growing realization that medical knowledge, originally codified for the purpose of computer-based consultations, may be utilized in additional ways that are medically relevant. Using the knowledge to teach medical students is perhaps foremost among these, and NEOMYCIN continues to focus on methods for augmenting clinical knowledge in order to facilitate its use in a tutorial setting. A particularly important aspect of this work is the insight that has been gained regarding the need to structure knowledge differently, and in more detail, when it is being used for different purposes (e.g., teaching as opposed to clinical decision making). It was this aspect of the

GUIDON research that led to the development of NEOMYCIN, which is an evolving computational model of medical diagnostic reasoning that we hope will enable us to better understand and teach diagnosis to students. An important additional realization is that these structuring methods are beneficial for improving the problem-solving performance of consultation programs, providing more detailed and abstract explanations to consultation users, and making knowledge bases easier to maintain.

As we move from technological development of explanation and student modeling capabilities, we are now collaborating closely with medical students and physicians to design an effective, useful tutoring program. Stanford Medical School faculty, such as Dr. Maffly, have shown considerable interest in this project. A research fellow associated with Maffly, Curt Kapsner, M.D., worked with the project for two years while at Stanford. The project also collaborates with a community of researchers focusing on medical education, funded by the Josiah Macy, Jr. Foundation.

C. Highlights of Research Progress

C.1 Accomplishments This Past Year

C.1.1 The NEOMYCIN Consultation Program

This year our physician and medical student exhaustively reviewed the knowledge base, filling in gaps and improving entries. Disease terminology and rule translation were modified, so that our medical collaborators now believe that the system is ready for use by students.

C.1.2 The GUIDON-WATCH Display Program

In the past few years, we have developed a complex interface for viewing NEOMYCIN's knowledge base and reasoning process. We call this program GUIDON-WATCH; it serves as the foundation for our teaching programs.

Major improvements were made this year to the display interface and underlying code. Interfacing a knowledge base with a bit-map display requires new programming methods. We standardized techniques for describing display structures (hierarchies, tables, and menus) and relating them to the various display operations (resizing text, highlighting subtrees, getting details, making hardcopy). We are teaching these coding methods to our new students, significantly improving their productivity and the clarity of the resulting program.

The display system for showing NEOMYCIN's reasoning was totally revamped to use multiple fonts for clarity, to allow uniform "buttoning" of any item to get additional information, and to simplify printout for a student. The simplification took several months in late 1985 and has been used by medical students in informal trials.

C.1.3 The GUIDON-DEBUG Instructional Program

During the past year, we developed a prototype version of GUIDON-DEBUG. Briefly, a student plays the role of a knowledge engineer, running cases and editing the knowledge base in order to find the missing or incorrect knowledge we have planted. The patient-specific model focuses his search, revealing what findings are not accounted for, suggesting what rules might be missing.

We provide extensive graphic facilities (GUIDON-WATCH), so the student needs to know little about the underlying implementation. Indeed, the idea is to make it easy to learn how NEOMYCIN works, rather than to hide what we are trying to teach. We have written about the theory of learning and novelty of our approach in publications cited elsewhere.

The idea of the *patient-specific model* window is the most significant technical advance in the past few months. This window shows the support for each disease hypothesis as a kind of proof. It quickly reveals gaps in the diagnostic explanation provided by the program, in a way that is hidden by the accumulating of evidence or scoring scheme used by most programs (including NEOMYCIN before we made this change). The idea of constructing this window comes from a mixture of sources: Patil's idea of a patient-specific model, Reggia's idea of "coverage" of disease hypotheses, Anderson's idea of showing a geometry proof as a graph, and Brown's idea that we should use graphics to help students visualize the reasoning process. Everyone, particularly our medical collaborators, has rallied around this conceptualization.

We have integrated our student modeling program, ODYSSEUS, into the GUIDON-DEBUG system. The modeling system is given a reordered set of data requests, with perhaps additional requests, which it parses with respect to NEOMYCIN's diagnostic strategy. That is, it attempts to find a strategic explanation for why each question was asked, abstracting coherent sequences from the list of data requests.

A new program relates this parse to the knowledge base, determining what difference to the knowledge base would produce the indicated ordering of questions. This serves both as an evaluation of a student's critique, as well as a very efficient knowledge acquisition device. It is gratifying to have this capability implemented, since it is the original idea that motivated the original design of NEOMYCIN.

An interesting capability was implemented that allows the user (student, physician, or knowledge engineer) to show the state of the consultation program at any time. In GUIDON-DEBUG you underline a question asked by NEOMYCIN and then display any window you wish. For example, if you say WHY, the program will answer the question as if you had asked it at that time during the consultation.

C.1.4 Dissemination of results

In the past year, we had a number of outstanding opportunities to present our work to other researchers. Dr. Clancey was an invited speaker at several European conferences, giving forty hours of lectures in Belgium, England, France, Germany, Holland, Italy, and Switzerland during this period. The above programs were demonstrated in five of these countries, as well as receiving strong attention at IJCAI in August in Los Angeles. (Demonstrations were given on Xerox Dandelions supplied by local representatives of Rank-Xerox. We are impressed by how easy it is to merely carry floppies to Europe and simply start up the program on the Xerox personal workstation.)

As the program becomes ready for demonstrating in tutorial mode, we are showing it to more people with a medical background. In January, Dr. Clancey gave a lecture and demonstration at a seminar run by the AAMC in San Diego (Information Management Technology). The program was also demonstrated at a Biomedical conference. We also entertained G. Octo Barnett for a morning, introducing him to the idea of intelligent tutoring systems. We are continuing to work with one of Dr. Maffly's medical students; Maffly is the main proponent of CAI in the Stanford Medical School. This student carries our work back to Maffly's lab, beginning the process of integrating our research with more traditional CAI practice.

C.2 Research in Progress

The following projects are active as of June 1986 (see also near-term plans listed in Section III.A):

1. Developing additional instructional programs based on NEOMYCIN;

2. Studying learning in the setting of debugging a knowledge base;

3. Re-implementing the explanation program to use the logic-encoding of the

metarules (stating this program in the same task/metarule language so that it might reason about its own explanations);

4. Generalizing our graphics package using object-oriented techniques;
5. Applying the student modeling program, ODYSSEUS, to knowledge acquisition; and
6. Preparing HERACLES, the generalization of NEOMYCIN, for use by other people.

D. Publications Since January 1985

1. (*) Clancey, W.J.: *Review of Sowa's "Conceptual Structures: Information Processing in Mind and Machine." Journal of Artificial Intelligence* 27(1):113-124, 1985.
2. (*) Clancey, W. J.: *Expert system design: Lessons from applications to education (abstract).* IEEE Spring Computer Conference, San Francisco, page 115.
3. (*) Thompson, T. and Clancey, W. J.: *A qualitative modeling shell for process diagnosis. IEEE Software* 3(2):6-15, March 1986.
4. (*) Clancey, W.J.:*The science and engineering of qualitative models.* Submitted to AAAI-86, March 1986.
5. (*) Clancey, W.J., Richer, M., Wilkins, D.C., Barnhouse, S., Kapsner, C., Leserman, D., Macias, J., Merchant, A., and Rodolitz, N.: *Guidon-Debug: The student as knowledge engineer.* Submitted to AAAI-86, April 1986.
6. (*) Clancey W.J.: *Qualitative Student Models.* **Annual Review of Computer Science.** Palo Alto: Annual Reviews, Report KSL-86-11, Computer Science Dept., May 1986.
7. (*) Clancey W.J.: *From GUIDON to NEOMYCIN and HERACLES in twenty short lessons: ONR Final Report, 1979-1985.* To appear in the *AI Magazine*, August 1986.

E. Funding Support

Contract Title: "A Family of Intelligent Tutoring Programs for Medical Diagnosis"
Principal Investigator: Bruce G. Buchanan, Prof. Computer Science, Research
Associate Investigator: William J. Clancey, Research Assoc. Computer Science
Agency: Josiah Macy, Jr. Foundation
Term: March 1985 to March 1988
Total award: $503,415 direct costs

Contract Title: "Computer-Based Tutors for Explaining and Managing the Process of Diagnostic Reasoning"
Principal Investigator: Bruce G. Buchanan, Prof. Computer Science, Research
Associate Investigator: William J. Clancey, Research Assoc. Computer Science
Agency: Office of Naval Research
ID number: N00014-85-K-0305
Total award: $510,311 total

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

A great deal of interest in GUIDON and NEOMYCIN has been shown by the medical and computer science communities. We are frequently asked to demonstrate these programs to Stanford visitors or at meetings in this country or abroad. Physicians have generally been enthusiastic about the potential of these programs and what they reveal about current approaches to computer-based medical decision making.

B. Sharing and Interaction with Other SUMEX-AIM Projects

GUIDON/NEOMYCIN retains strong contact with the ONCOCIN project, as both are siblings of the MYCIN parent. These projects share programming expertise and utility routines. In addition, the central SUMEX development group acts as an important clearing house for solving problems and distributing new methods.

C. Critique of Resource Management

The SUMEX staff has been extremely helpful in maintaining connections between Xerox D-machines and SUMEX. The SUMEX group's aggressive program of converting computing resources to the workstation environment has been an outstanding success. Computing for the average programmer at any hour of the day is better than anyone experienced ten years ago, except for early morning on an unloaded system. With the addition of bit-map graphics within the programming environment and sufficient printers, we find that average preparation time for a figure to be inserted in a paper has dropped from several hours to one hour or even just a few minutes.

III. RESEARCH PLANS

A. Project Goals and Plans

Research over the next year will continue on several fronts, including one or more prototype instructional programs.

1. Use of GUIDON-DEBUG by medical students to empirically develop the interface and teaching scenario. In this setting, we will have physician teachers watch and advise as students use the program, using this opportunity to formulate their teaching strategies.

2. Incremental improvements to GUIDON-DEBUG, including work on editing the patient-specific model, evaluating student annotations of the transcript, and probably some kind of advising or coaching capability.

3. Starting in the fall, we will begin implementing a new explanation program that will produce clear text, use a student model, and be integrated with the display package.

4. A new project idea is to have online text, such as from the book Medicine, by Fishman, et al., tied to the knowledge base. A replica of the book, with fonts, will be displayed on the screen. Selecting text will display the relevant portion of the knowledge base. This will be a good entry point for students, as well as providing an interesting introduction to the idea of a knowledge base for demonstrations to physicians.

5. Another new idea which would be easy and useful to do, is to have a means for recording interactions with the program. These sessions would then be

edited for redisplay as automatic demonstrations or as introductions for new students. Text, questions, and the opportunity to try certain features will involve the viewer and test his understanding. We call this system GUIDON-SCRIPTS. This capability to replay display sequences would also be invoked automatically by explanation and teaching rules.

B. Long term plans: the GUIDON2 Family of Instructional Programs

In last year's annual report, we described our research program for 1985-88. Here, we provide a summary of the most important projects that may be completed in the next few years.

GUIDON-MANAGE In this system the student solves a problem by telling NEOMYCIN what task to do at each step. Essentially, the student provides the strategy and the program supplies the tactics (meta-rules) and domain knowledge to carry out the strategy. The program will in general carry through tasks in a logical way, for example, proceeding to test a hypothesis completely, and not "breaking" on low-level tasks that mainly test domain knowledge rather than strategy. The program will not pursue new hypotheses automatically. However, the student will always see what questions a task caused the program to request, as well as how the differential changes. This activity leads the student to observe what a strategy entails, helping him become a better observer of his own behavior. Here he shows that he knows the structural vocabulary that makes a strategy appropriate.

GUIDON-APPRENTICE This is a variant of NEOMYCIN in which the program stops during a consultation and asks the student to propose the next data request(s). The student is asked to indicate the task and focus he has in mind, plus the differential he is operating upon. The program compares this proposal to what NEOMYCIN would do. In this activity we descend to the domain level and require the student to instantiate a strategy appropriately. Ultimately, such a program will use a *learning model* that anticipates what the student is ready to learn next and how he should be challenged. Early versions can simply use built-in breakpoints supplied by an expert teacher. In the future, programs will develop their own curricula from a case library.

GUIDON-SOLVE This is the complete tutorial system. The student carries through diagnosis completely, while a student modeling program attempts to track what he is doing and a coach interrupts to offer advice. Here annotation, comparison, debugging, and explanation are all integrated to illustrate to the student how his solution is non-optimal. For example, the student might be asked to annotate his solution after he is done; this will point out strategic gaps in his awareness and provide a basis for critique and improvement. A "curriculum" based on frequent student faults and important things to learn will drive the interaction. In this activity, the student is on his own. Faced with the proverbial "blank screen," he must exercise his diagnostic procedure from start to finish.

Now that we are producing instructional programs from NEOMYCIN which can be tested, we will be emphasizing empirical work with students in all of our future research.

C. Requirements for Continued SUMEX Use

The D-machine's large address space permits development of the large programs that complex computer-aided instruction requires. Graphics enable us to develop new methods for presenting material to naive users. We use the D-machine as a reliable, constant "load-average" machine, for running experiments with physicians and students. The development of GUIDON2 on the D-machine demonstrates the feasibility of running intelligent consultation or tutoring systems on small, affordable machines in physicians' offices, schools, and other remote sites.